files for the distributed storage system, are combined with one or more distributed object storage managers ("DOSMs"). A network couples the DOSMs to the intelligent storage nodes. The DOSMs manage requests from clients, such as content servers and end-user computers, to download files from a storage center.

The distributed storage system is arranged to support uninterrupted delivery of files in the event of a failure in an intelligent storage node. In one embodiment, the failover architecture includes storing, for each file, a duplicate file in a different intelligent storage node. In the event the distributed storage system enters a failover condition, the DOSM determines, for files stored in the failed intelligent storage node, a location (*i.e.*, intelligent storage node) for the duplicate files. In one embodiment, to determine the location of the intelligent storage node storing a duplicate file, the DOSMs store a map. The map provides a correspondence between a network address of the failed intelligent storage node and a network address of the intelligent storage node that stores the duplicate file.

In one embodiment, the intelligent storage nodes, which store duplicate files, are located in different storage centers (*i.e.*, first and second storage centers). The storage centers are located in different geographic areas (*e.g.*, west coast of United States and east coast of United States).

The underlying file system for the files, referred to as a virtual file system

("VFS"), is separate from the storage cluster. The VFS includes one or more distributed directory managers ("DDMs") and one or more directories. The directories store the file system information, and the DDMs manage requests from clients for file system information. The file system contents are duplicated in more than one directory. In the event of a failover condition, the DDMs access another directory (*e.g.*, computer) to obtain the file system contents.

Rejection of the Claims Under 35 U.S.C. § 102

In the Office Action dated November 19, 2002, claims 1, 2, 4-7, 11, 12, 14-17, 21, 22, 24-27 were rejected under 35 U.S.C. § 103 as being unpatentable over US Patent 5,692,155, issued to *Iskiyan et al.*, in view of US Patent 5,870,537, issued to *Kern et al.* Claims 3, 13 and 23 were rejected under 35 USC § 103 as being unpatentable over *Iskiyan et al.* in view of *Kern et al.*, and in further view of *Mogul.*

Overview of Cited References:

Iskiyan et al. disclose a synchronous dual copy data storage system. (Col. 3, lines 22-25). Specifically, the synchronous dual copy data storage system consists of a first and second duplex pair. Each duplex pair has a primary storage device and a secondary storage device. A data storage system, with a primary storage device and a secondary storage device, is shown in Fig. 1. The primary site includes a host processor, storage controller, and attached primary disk attached storage device "DASD." The secondary

site includes a secondary processor, connected to a secondary controller, which in turn, is connected to a DASD. The storage controller directs copies of the records and record updates to the secondary DASDs of the duplex pairs.

Kern et al. also discloses a duplex pair that includes a primary data storage device and secondary data storage device. The secondary data storage device is swapped with the primary storage device by terminating the remote copy duplex pair, establishing an opposite direction remote copy duplex pair such that the secondary data storage device is a primary storage device and the primary data storage device is the shadowing device. (Abstract).

Claims 1, 11 and 21:

A. *Iskiyan et al.* And *Kern et al.* Disclose A Duplex Pair In A Storage System.

Amended claim 1 recites, in part:

providing a plurality of a plurality of Distributed Object Storage Managers "DOSMs" for receiving requests for files;

providing at least three intelligent storage nodes accessible to said DOSMs over a network;

storing at least one file in a first intelligent storage node accessed via a DOSM over said network;

storing a duplicate of said file in a second intelligent storage node accessed via said network;

As claimed, amended claim 1 includes a plurality of distributed object storage managers that access a plurality of intelligent storage nodes over a network. The DOSMs receive file requests. The intelligent storage nodes store a file in two of at least three intelligent storage nodes in the system.

Iskiyan et al. And *Kern et al.* both disclose a duplex pair, such that a primary file system stores a file and the secondary file system stores a duplicate of the file. Thus, in the *Iskiyan et al.* and *Kern et al.* systems, the same two devices are always used to store the file and the copy of the file. In contrast, the claimed invention claims at least three intelligent storage nodes. As such, the locations of the file and the duplicate of the file are not dedicated to two specific devices (*i.e.*, a primary and a secondary storage device). Accordingly, *Iskiyan et al.* and *Kern et al.* do not disclose or teach a storage system that has at least three nodes for storage of a file and a duplicate file on two intelligent storage nodes.

B. The Storage Controllers Of *Iskiyan et al.* Do Not Determine A Location For A File If A Failover Condition Occurs.

Amended claim 1 sets forth:

determining, at said DOSM, a location for said file in one of said intelligent storage nodes;

identifying said second storage node as said location for said file;

Iskiyan et al. and *Kern et al.* designate a primary storage device and a secondary storage device. As such, if a failure occurs, the system selects the secondary storage device. In contrast, with the storage architecture of the claimed invention, the system must determine a location of the duplicate copy. Specifically, as claimed in claim 1, a DOSM determines the location for a duplicate file in one of the intelligent storage nodes. *Iskiyan et al.* and *Kern et al.* do not determine the location for a duplicate file in one of the intelligent storage nodes, and therefore the cited references do not render the claimed invention obvious.

C. The Claimed Invention Supports A Failover System For A Storage Architecture That Has More Than Two Storage Nodes.

The storage architecture of the claimed invention permits storing a large amount of files across multiple intelligent storage nodes. In addition, the storage architecture does not require dedicating a single storage node for primary file storage and another single storage node for duplicate storage. Moreover, the cited references do not suggest a means for executing a failover condition in a storage system that has at least three storage nodes accessed by a plurality of managers. For these additional reasons, the cited references do not render the claimed invention obvious.

CONCLUSION

In view of the foregoing, it is submitted that the claims are in condition for allowance. Reconsideration of the rejections and objections is requested. Allowance is earnestly solicited at the earliest possible date.

Respectfully submitted,

STATTLER JOHANSEN & ADELI LLP

Dated: 5/19/03

John Stattler
Reg. No. 36,285

Stattler, Johansen & Adeli LLP
PO Box 51860
Palo Alto, CA 94303-0728
Phone: (650) 752-0990 ext.100
Fax: (650) 752-0995

The Amended Claims

The following pages provide the amended claims with the amendments marked with deleted material in [brackets] and new material underlined to show the changes made.

1. (Once Amended) A method for configuring a distributed storage system for a failover condition, said method comprising the steps of:

providing a plurality of distributed object storage managers "DOSMs" for receiving requests for files;

providing at least three intelligent storage nodes accessible to said DOSMs over a network;

storing at least one file in a first intelligent storage node accessed via a DOSM over said network;

storing a duplicate of said file in a second intelligent storage node accessed via said network;

entering a failover condition to cease use of said first intelligent storage node;

determining, at said DOSM, a location for said file in one of said [second] intelligent storage nodes; [and]

identifying said second storage node as said location for said file; and

accessing, via said network, said file stored in said second intelligent storage node in response to a subsequent file request[;].

8. (Once Amended) The method as set forth in claim 7, further comprising the step of searching for said file using a point-to-point protocol between [a] said distributed object storage manager (DOSM) and said second [an] intelligent storage node.

11. (Once Amended) A distributed storage system comprising:

a network;

at least three intelligent storage nodes;

a first intelligent storage node, accessed via said network, for storing at least one file;

a second intelligent storage node, accessed via said network, for storing a duplicate of said file; and

[at least one] a plurality of distributed object storage managers (DOSMs), one of said DOSMs for accessing, via said network, said file stored in said first intelligent storage node in response to a file request, for entering a failover condition to cease access to said file in said first intelligent storage node, for determining a location for said file in one of said [second] intelligent storage nodes, for identifying said second storage node as said location for said file, and for accessing, via said network, said file stored in said second intelligent storage node in response to a subsequent file request.